

Package ‘EQRN’

May 7, 2026

Type Package

Title Extreme Quantile Regression Neural Networks for Risk Forecasting

Version 0.1.2

Description This framework enables forecasting and extrapolating measures of conditional risk (e.g. of extreme or unprecedented events), including quantiles and exceedance probabilities, using extreme value statistics and flexible neural network architectures. It allows for capturing complex multivariate dependencies, including dependencies between observations, such as sequential dependence (time-series). The methodology was introduced in Pasche and Engelke (2024) <[doi:10.1214/24-AOAS1907](https://doi.org/10.1214/24-AOAS1907)> (also available in preprint: Pasche and Engelke (2022) <[doi:10.48550/arXiv.2208.07590](https://doi.org/10.48550/arXiv.2208.07590)>).

License GPL (>= 3)

Encoding UTF-8

Imports coro (>= 1.0.2), doFuture, evd, foreach, future, ismev, magrittr, stats, torch, utils

RoxygenNote 7.3.2

URL <https://github.com/opasche/EQRN>, <https://opasche.github.io/EQRN/>

BugReports <https://github.com/opasche/EQRN/issues>

NeedsCompilation no

Author Olivier C. Pasche [aut, cre, cph] (ORCID: <<https://orcid.org/0000-0002-1202-9199>>)

Maintainer Olivier C. Pasche <olivier_pasche@alumni.epfl.ch>

Repository CRAN

Date/Publication 2025-11-21 13:10:02 UTC

Contents

backend_is_installed	3
check_directory	4
compute_EQRN_GPDLoss	4
compute_EQRN_seq_GPDLoss	5

default_device	6
end_doFuture_strategy	7
EQRN_excess_probability	7
EQRN_excess_probability_seq	8
EQRN_fit	9
EQRN_fit_restart	12
EQRN_fit_seq	13
EQRN_load	15
EQRN_predict	16
EQRN_predict_params	16
EQRN_predict_params_seq	17
EQRN_predict_seq	18
EQRN_save	19
excess_probability	20
excess_probability.EQRN_iid	20
excess_probability.EQRN_seq	21
FC_GPD_net	22
FC_GPD_SNN	23
fit_GPD_unconditional	24
get_doFuture_operator	25
get_excesses	25
GPD_excess_probability	26
GPD_quantiles	27
install_backend	28
lagged_features	28
last_elem	29
loss_GPD	29
loss_GPD_tensor	30
make_folds	31
mean_absolute_error	31
mean_squared_error	32
mts_dataset	33
multilevel_exceedance_proba_error	34
multilevel_MAE	35
multilevel_MSE	36
multilevel_pred_bias	37
multilevel_prop_below	38
multilevel_q_loss	39
multilevel_q_pred_error	40
multilevel_resid_var	41
multilevel_R_squared	42
perform_scaling	43
predict.EQRN_iid	43
predict.EQRN_seq	44
predict.QRN_seq	45
prediction_bias	46
prediction_residual_variance	46
predict_GPD_semiconditional	47

predict_unconditional_quantiles	48
process_features	49
proportion_below	50
QRNN_RNN_net	50
QRN_fit_multiple	51
QRN_seq_fit	52
QRN_seq_predict	54
QRN_seq_predict_foldwise	55
QRN_seq_predict_foldwise_sep	56
quantile_exceedance_proba_error	57
quantile_loss	58
quantile_loss_tensor	59
quantile_prediction_error	59
Recurrent_GPD_net	60
roundm	61
R_squared	62
safe_save_rds	62
semiconditional_train_valid_GPD_loss	63
Separated_GPD_SNN	64
set_doFuture_strategy	65
square_loss	66
unconditional_train_valid_GPD_loss	66
vec2mat	67
vector_insert	68
Index	69

backend_is_installed *Check if Torch Backend Libraries are Installed*

Description

Check if Torch Backend Libraries are Installed

Usage

```
backend_is_installed(...)
```

Arguments

... Optional parameters passed to `torch::torch_is_installed()`.

Value

Boolean indicating whether the LibTorch and LibLantern backend libraries are installed.

Examples

```
if(backend_is_installed()){
  cat("torch's LibTorch and LibLantern backend libraries are installed!\n")
}
```

check_directory	<i>Check directory existence</i>
-----------------	----------------------------------

Description

Checks if the desired directory exists. If not, the desired directory is created.

Usage

```
check_directory(dir_name, recursive = TRUE, no_warning = FALSE)
```

Arguments

dir_name	Path to the desired directory, as a string.
recursive	Should elements of the path other than the last be created? If TRUE, behaves like the Unix command <code>mkdir -p</code> .
no_warning	Whether to cancel the warning issued if a directory is created (bool).

Value

No return value.

Examples

```
check_directory("./some_folder/my_new_folder")
```

compute_EQRN_GPDLoss	<i>Generalized Pareto likelihood loss of a EQRN_iid predictor</i>
----------------------	---

Description

Generalized Pareto likelihood loss of a EQRN_iid predictor

Usage

```
compute_EQRN_GPDLoss(
  fit_eqrn,
  X,
  y,
  intermediate_quantiles = NULL,
  interm_lvl = fit_eqrn$interm_lvl,
  device = default_device()
)
```

Arguments

fit_eqrn	Fitted "EQRN_iid" object.
X	Matrix of covariates.
y	Response variable vector.
intermediate_quantiles	Vector of intermediate conditional quantiles at level fit_eqrn\$interm_lvl.
interm_lvl	Optional, checks that interm_lvl == fit_eqrn\$interm_lvl.
device	(optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .

Value

Negative GPD log likelihood of the conditional EQRN predicted parameters over the response exceedances over the intermediate quantiles.

```
compute_EQRN_seq_GPDLoss
```

Generalized Pareto likelihood loss of a EQRN_seq predictor

Description

Generalized Pareto likelihood loss of a EQRN_seq predictor

Usage

```
compute_EQRN_seq_GPDLoss(
  fit_eqrn,
  X,
  Y,
  intermediate_quantiles = NULL,
  interm_lvl = fit_eqrn$interm_lvl,
  seq_len = fit_eqrn$seq_len,
  device = default_device()
)
```

Arguments

<code>fit_eqrn</code>	Fitted "EQRN_seq" object.
<code>X</code>	Matrix of covariates.
<code>Y</code>	Response variable vector corresponding to the rows of <code>X</code> .
<code>intermediate_quantiles</code>	Vector of intermediate conditional quantiles at level <code>fit_eqrn\$interm_lvl</code> .
<code>interm_lvl</code>	Optional, checks that <code>interm_lvl == fit_eqrn\$interm_lvl</code> .
<code>seq_len</code>	Data sequence length (i.e. number of past observations) used to predict each response quantile. By default, the training <code>fit_eqrn\$seq_len</code> is used.
<code>device</code>	(optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .

Value

Negative GPD log likelihood of the conditional EQRN predicted parameters over the response exceedances over the intermediate quantiles.

<code>default_device</code>	<i>Default torch device</i>
-----------------------------	-----------------------------

Description

Default torch device

Usage

```
default_device()
```

Value

Returns `torch::torch_device("cuda")` if `torch::cuda_is_available()`, or `torch::torch_device("cpu")` otherwise.

Examples

```
if(backend_is_installed()){
  device <- default_device()
  print(device)
}
```

end_doFuture_strategy *End the currently set doFuture strategy*

Description

Resets the default strategy using `future::plan("default")`.

Usage

```
end_doFuture_strategy()
```

Value

No return value.

Examples

```
`%fun%` <- set_doFuture_strategy("multisession", n_workers=3)
# perform foreach::foreach loop using the %fun% operator
end_doFuture_strategy()
```

EQRN_excess_probability

Tail excess probability prediction using an EQRN_iid object

Description

Tail excess probability prediction using an `EQRN_iid` object

Usage

```
EQRN_excess_probability(
  val,
  fit_eqrn,
  X,
  intermediate_quantiles,
  interm_lvl = fit_eqrn$interm_lvl,
  body_proba = "default",
  proba_type = c("excess", "cdf"),
  device = default_device()
)
```

Arguments

<code>val</code>	Quantile value(s) used to estimate the conditional excess probability or cdf.
<code>fit_eqrn</code>	Fitted "EQRN_iid" object.
<code>X</code>	Matrix of covariates to predict the corresponding response's conditional excess probabilities.
<code>intermediate_quantiles</code>	Vector of intermediate conditional quantiles at level <code>fit_eqrn\$interm_lvl</code> .
<code>interm_lvl</code>	Optional, checks that <code>interm_lvl == fit_eqrn\$interm_lvl</code> .
<code>body_proba</code>	Value to use when the predicted conditional probability is below <code>interm_lvl</code> (in which case it cannot be precisely assessed by the model). If "default" is given (the default), <code>paste0(">", 1-interm_lvl)</code> is used if <code>proba_type=="excess"</code> , and <code>paste0("<", interm_lvl)</code> is used if <code>proba_type=="cdf"</code> .
<code>proba_type</code>	Whether to return the "excess" probability over <code>val</code> (default) or the "cdf" at <code>val</code> .
<code>device</code>	(optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .

Value

Vector of probabilities (and possibly a few `body_proba` values if `val` is not large enough) of length `nrow(X)`.

EQRN_excess_probability_seq

Tail excess probability prediction using an EQRN_seq object

Description

Tail excess probability prediction using an EQRN_seq object

Usage

```
EQRN_excess_probability_seq(
  val,
  fit_eqrn,
  X,
  Y,
  intermediate_quantiles,
  interm_lvl = fit_eqrn$interm_lvl,
  crop_predictions = FALSE,
  body_proba = "default",
  proba_type = c("excess", "cdf"),
  seq_len = fit_eqrn$seq_len,
  device = default_device()
)
```

Arguments

val	Quantile value(s) used to estimate the conditional excess probability or cdf.
fit_eqrn	Fitted "EQRN_seq" object.
X	Matrix of covariates to predict the response's conditional excess probabilities.
Y	Response variable vector corresponding to the rows of X.
intermediate_quantiles	Vector of intermediate conditional quantiles at level <code>fit_eqrn\$interm_lvl</code> .
interm_lvl	Optional, checks that <code>interm_lvl == fit_eqrn\$interm_lvl</code> .
crop_predictions	Whether to crop out the first <code>seq_len</code> observations (which are NA) from the returned vector
body_proba	Value to use when the predicted conditional probability is below <code>interm_lvl</code> (in which case it cannot be precisely assessed by the model). If "default" is given (the default), <code>paste0(">", 1-interm_lvl)</code> is used if <code>proba_type=="excess"</code> , and <code>paste0("<", interm_lvl)</code> is used if <code>proba_type=="cdf"</code> .
proba_type	Whether to return the "excess" probability over <code>val</code> (default) or the "cdf" at <code>val</code> .
seq_len	Data sequence length (i.e. number of past observations) used to predict each response quantile. By default, the training <code>fit_eqrn\$seq_len</code> is used.
device	(optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .

Value

Vector of probabilities (and possibly a few `body_proba` values if `val` is not large enough) of length `nrow(X)` (or `nrow(X)-seq_len` if `crop_predictions`).

EQRN_fit

EQRN fit function for independent data

Description

Use the `EQRN_fit_restart()` wrapper instead, with `data_type="iid"`, for better stability using fitting restart.

Usage

```
EQRN_fit(
  X,
  y,
  intermediate_quantiles,
  interm_lvl,
  shape_fixed = FALSE,
  net_structure = c(5, 3, 3),
  hidden_fct = torch::nnf_sigmoid,
```

```

p_drop = 0,
intermediate_q_feature = TRUE,
learning_rate = 1e-04,
L2_pen = 0,
shape_penalty = 0,
scale_features = TRUE,
n_epochs = 500,
batch_size = 256,
X_valid = NULL,
y_valid = NULL,
quant_valid = NULL,
lr_decay = 1,
patience_decay = n_epochs,
min_lr = 0,
patience_stop = n_epochs,
tol = 1e-06,
orthogonal_gpd = TRUE,
patience_lag = 1,
optim_met = "adam",
seed = NULL,
verbose = 2,
device = default_device()
)

```

Arguments

X	Matrix of covariates, for training.
y	Response variable vector to model the extreme conditional quantile of, for training.
intermediate_quantiles	Vector of intermediate conditional quantiles at level <code>interm_lvl</code> .
interm_lvl	Probability level for the intermediate quantiles <code>intermediate_quantiles</code> .
shape_fixed	Whether the shape estimate depends on the covariates or not (bool).
net_structure	Vector of integers whose length determines the number of layers in the neural network and entries the number of neurons in each corresponding successive layer. If <code>hidden_fct=="SSNN"</code> , should instead be a named list with "scale" and "shape" vectors for the two respective sub-networks. Can also be a torch::nn_module network with correct input and output dimensions, which overrides the <code>hidden_fct</code> , <code>shape_fixed</code> and <code>p_drop</code> arguments.
hidden_fct	Activation function for the hidden layers. Can be either a callable function (preferably from the torch library), or one of the the strings "SNN", "SSNN" for self normalizing networks (with common or separated networks for the scale and shape estimates, respectively). In the latter cases, <code>shape_fixed</code> has no effect.
p_drop	Probability parameter for dropout before each hidden layer for regularization during training. alpha-dropout is used with SNNs.

<code>intermediate_q_feature</code>	Whether to use the <code>intermediate_quantiles</code> as an additional covariate, by appending it to the X matrix (bool).
<code>learning_rate</code>	Initial learning rate for the optimizer during training of the neural network.
<code>L2_pen</code>	L2 weight penalty parameter for regularization during training.
<code>shape_penalty</code>	Penalty parameter for the shape estimate, to potentially regularize its variation from the fixed prior estimate.
<code>scale_features</code>	Whether to rescale each input covariates to zero mean and unit variance before applying the network (recommended).
<code>n_epochs</code>	Number of training epochs.
<code>batch_size</code>	Batch size used during training.
<code>X_valid</code>	Covariates in a validation set, or NULL. Used for monitoring validation loss during training, enabling learning-rate decay and early stopping.
<code>y_valid</code>	Response variable in a validation set, or NULL. Used for monitoring validation loss during training, enabling learning-rate decay and early stopping.
<code>quant_valid</code>	Intermediate conditional quantiles at level <code>interm_lvl</code> in a validation set, or NULL. Used for monitoring validation loss during training, enabling learning-rate decay and early stopping.
<code>lr_decay</code>	Learning rate decay factor.
<code>patience_decay</code>	Number of epochs of non-improving validation loss before a learning-rate decay is performed.
<code>min_lr</code>	Minimum learning rate, under which no more decay is performed.
<code>patience_stop</code>	Number of epochs of non-improving validation loss before early stopping is performed.
<code>tol</code>	Tolerance for stopping training, in case of no significant training loss improvements.
<code>orthogonal_gpd</code>	Whether to use the orthogonal reparametrization of the estimated GPD parameters (recommended).
<code>patience_lag</code>	The validation loss is considered to be non-improving if it is larger than on any of the previous <code>patience_lag</code> epochs.
<code>optim_met</code>	DEPRECATED. Optimization algorithm to use during training. "adam" is the default.
<code>seed</code>	Integer random seed for reproducibility in network weight initialization.
<code>verbose</code>	Amount of information printed during training (0:nothing, 1:most important, 2:everything).
<code>device</code>	(optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .

Value

An EQRN object of classes `c("EQRN_iid", "EQRN")`, containing the fitted network, as well as all the relevant information for its usage in other functions.

EQRN_fit_restart *Wrapper for fitting EQRN with restart for stability*

Description

Wrapper for fitting EQRN with restart for stability

Usage

```
EQRN_fit_restart(
  X,
  y,
  intermediate_quantiles,
  interm_lvl,
  number_fits = 3,
  ...,
  seed = NULL,
  data_type = c("iid", "seq")
)
```

Arguments

<code>X</code>	Matrix of covariates, for training.
<code>y</code>	Response variable vector to model the extreme conditional quantile of, for training.
<code>intermediate_quantiles</code>	Vector of intermediate conditional quantiles at level <code>interm_lvl</code> .
<code>interm_lvl</code>	Probability level for the intermediate quantiles <code>intermediate_quantiles</code> .
<code>number_fits</code>	Number of restarts.
<code>...</code>	Other parameters given to either EQRN_fit() or EQRN_fit_seq() , depending on the <code>data_type</code> .
<code>seed</code>	Integer random seed for reproducibility in network weight initialization.
<code>data_type</code>	Type of data dependence, must be one of "iid" (for iid observations) or "seq" (for sequentially dependent observations).

Value

An EQRN object of classes `c("EQRN_iid", "EQRN")`, if `data_type=="iid"`, or `c("EQRN_seq", "EQRN")`, if `data_type=="seq"`, containing the fitted network, as well as all the relevant information for its usage in other functions.

`EQRN_fit_seq`*EQRN fit function for sequential and time series data*

Description

Use the `EQRN_fit_restart()` wrapper instead, with `data_type="seq"`, for better stability using fitting restart.

Usage

```
EQRN_fit_seq(  
  X,  
  y,  
  intermediate_quantiles,  
  interm_lvl,  
  shape_fixed = FALSE,  
  hidden_size = 10,  
  num_layers = 1,  
  rnn_type = c("lstm", "gru"),  
  p_drop = 0,  
  intermediate_q_feature = TRUE,  
  learning_rate = 1e-04,  
  L2_pen = 0,  
  seq_len = 10,  
  shape_penalty = 0,  
  scale_features = TRUE,  
  n_epochs = 500,  
  batch_size = 256,  
  X_valid = NULL,  
  y_valid = NULL,  
  quant_valid = NULL,  
  lr_decay = 1,  
  patience_decay = n_epochs,  
  min_lr = 0,  
  patience_stop = n_epochs,  
  tol = 1e-05,  
  orthogonal_gpd = TRUE,  
  patience_lag = 1,  
  fold_separation = NULL,  
  optim_met = "adam",  
  seed = NULL,  
  verbose = 2,  
  device = default_device()  
)
```

Arguments

`X` Matrix of covariates, for training. Entries must be in sequential order.

<code>y</code>	Response variable vector to model the extreme conditional quantile of, for training. Entries must be in sequential order.
<code>intermediate_quantiles</code>	Vector of intermediate conditional quantiles at level <code>interm_lvl</code> .
<code>interm_lvl</code>	Probability level for the intermediate quantiles <code>intermediate_quantiles</code> .
<code>shape_fixed</code>	Whether the shape estimate depends on the covariates or not (bool).
<code>hidden_size</code>	Dimension of the hidden latent state variables in the recurrent network.
<code>num_layers</code>	Number of recurrent layers.
<code>rnn_type</code>	Type of recurrent architecture, can be one of "lstm" (default) or "gru".
<code>p_drop</code>	Probability parameter for dropout before each hidden layer for regularization during training.
<code>intermediate_q_feature</code>	Whether to use the <code>intermediate_quantiles</code> as an additional covariate, by appending it to the X matrix (bool).
<code>learning_rate</code>	Initial learning rate for the optimizer during training of the neural network.
<code>L2_pen</code>	L2 weight penalty parameter for regularization during training.
<code>seq_len</code>	Data sequence length (i.e. number of past observations) used during training to predict each response quantile.
<code>shape_penalty</code>	Penalty parameter for the shape estimate, to potentially regularize its variation from the fixed prior estimate.
<code>scale_features</code>	Whether to rescale each input covariates to zero mean and unit covariance before applying the network (recommended).
<code>n_epochs</code>	Number of training epochs.
<code>batch_size</code>	Batch size used during training.
<code>X_valid</code>	Covariates in a validation set, or NULL. Entries must be in sequential order. Used for monitoring validation loss during training, enabling learning-rate decay and early stopping.
<code>y_valid</code>	Response variable in a validation set, or NULL. Entries must be in sequential order. Used for monitoring validation loss during training, enabling learning-rate decay and early stopping.
<code>quant_valid</code>	Intermediate conditional quantiles at level <code>interm_lvl</code> in a validation set, or NULL. Used for monitoring validation loss during training, enabling learning-rate decay and early stopping.
<code>lr_decay</code>	Learning rate decay factor.
<code>patience_decay</code>	Number of epochs of non-improving validation loss before a learning-rate decay is performed.
<code>min_lr</code>	Minimum learning rate, under which no more decay is performed.
<code>patience_stop</code>	Number of epochs of non-improving validation loss before early stopping is performed.
<code>tol</code>	Tolerance for stopping training, in case of no significant training loss improvements.

orthogonal_gpd	Whether to use the orthogonal reparametrization of the estimated GPD parameters (recommended).
patience_lag	The validation loss is considered to be non-improving if it is larger than on any of the previous patience_lag epochs.
fold_separation	Index of fold separation or sequential discontinuity in the data.
optim_met	DEPRECATED. Optimization algorithm to use during training. "adam" is the default.
seed	Integer random seed for reproducibility in network weight initialization.
verbose	Amount of information printed during training (0:nothing, 1:most important, 2:everything).
device	(optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .

Value

An EQRN object of classes `c("EQRN_seq", "EQRN")`, containing the fitted network, as well as all the relevant information for its usage in other functions.

EQRN_load	<i>Load an EQRN object from disc</i>
-----------	--------------------------------------

Description

Loads in memory an "EQRN" object that has previously been saved on disc using `EQRN_save()`.

Usage

```
EQRN_load(path, name = NULL, device = default_device(), ...)
```

Arguments

path	Path to the save location as a string.
name	String name of the save. If NULL (default), assumes the save name has been given implicitly in the path.
device	(optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .
...	DEPRECATED. Used for back-compatibility.

Value

The loaded "EQRN" model.

EQRN_predict	<i>Predict function for an EQRN_iid fitted object</i>
--------------	---

Description

Predict function for an EQRN_iid fitted object

Usage

```
EQRN_predict(
  fit_eqrn,
  X,
  prob_lvls_predict,
  intermediate_quantiles,
  interm_lvl = fit_eqrn$interm_lvl,
  device = default_device()
)
```

Arguments

fit_eqrn	Fitted "EQRN_iid" object.
X	Matrix of covariates to predict the corresponding response's conditional quantiles.
prob_lvls_predict	Vector of probability levels at which to predict the conditional quantiles.
intermediate_quantiles	Vector of intermediate conditional quantiles at level fit_eqrn\$interm_lvl.
interm_lvl	Optional, checks that interm_lvl == fit_eqrn\$interm_lvl.
device	(optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .

Value

Matrix of size `nrow(X)` times `prob_lvls_predict` containing the conditional quantile estimates of the response associated to each covariate observation at each probability level. Simplifies to a vector if `length(prob_lvls_predict)==1`.

EQRN_predict_params	<i>GPD parameters prediction function for an EQRN_iid fitted object</i>
---------------------	---

Description

GPD parameters prediction function for an EQRN_iid fitted object

Usage

```
EQRN_predict_params(
  fit_eqrn,
  X,
  intermediate_quantiles = NULL,
  return_parametrization = c("classical", "orthogonal"),
  interm_lvl = fit_eqrn$interm_lvl,
  device = default_device()
)
```

Arguments

`fit_eqrn` Fitted "EQRN_iid" object.

`X` Matrix of covariates to predict conditional GPD parameters.

`intermediate_quantiles` Vector of intermediate conditional quantiles at level `fit_eqrn$interm_lvl`.

`return_parametrization` Which parametrization to return the parameters in, either "classical" or "orthogonal".

`interm_lvl` Optional, checks that `interm_lvl == fit_eqrn$interm_lvl`.

`device` (optional) A `torch::torch_device()`. Defaults to `default_device()`.

Value

Named list containing: "scales" and "shapes" as numerical vectors of length `nrow(X)`.

EQRN_predict_params_seq

GPD parameters prediction function for an EQRN_seq fitted object

Description

GPD parameters prediction function for an EQRN_seq fitted object

Usage

```
EQRN_predict_params_seq(
  fit_eqrn,
  X,
  Y,
  intermediate_quantiles = NULL,
  return_parametrization = c("classical", "orthogonal"),
  interm_lvl = fit_eqrn$interm_lvl,
  seq_len = fit_eqrn$seq_len,
  device = default_device()
)
```

Arguments

fit_eqrn	Fitted "EQRN_seq" object.
X	Matrix of covariates to predict conditional GPD parameters.
Y	Response variable vector corresponding to the rows of X.
intermediate_quantiles	Vector of intermediate conditional quantiles at level fit_eqrn\$interm_lvl.
return_parametrization	Which parametrization to return the parameters in, either "classical" or "orthogonal".
interm_lvl	Optional, checks that interm_lvl == fit_eqrn\$interm_lvl.
seq_len	Data sequence length (i.e. number of past observations) used to predict each response quantile. By default, the training fit_eqrn\$seq_len is used.
device	(optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .

Value

Named list containing: "scales" and "shapes" as numerical vectors of length nrow(X), and the seq_len used.

EQRN_predict_seq	<i>Predict function for an EQRN_seq fitted object</i>
------------------	---

Description

Predict function for an EQRN_seq fitted object

Usage

```
EQRN_predict_seq(
  fit_eqrn,
  X,
  Y,
  prob_lvls_predict,
  intermediate_quantiles,
  interm_lvl,
  crop_predictions = FALSE,
  seq_len = fit_eqrn$seq_len,
  device = default_device()
)
```

Arguments

fit_eqrn	Fitted "EQRN_seq" object.
X	Matrix of covariates to predict the corresponding response's conditional quantiles.

Y	Response variable vector corresponding to the rows of X.
prob_lvls_predict	Vector of probability levels at which to predict the conditional quantiles.
intermediate_quantiles	Vector of intermediate conditional quantiles at level <code>fit_eqrn\$interm_lvl</code> .
interm_lvl	Optional, checks that <code>interm_lvl == fit_eqrn\$interm_lvl</code> .
crop_predictions	Whether to crop out the first <code>seq_len</code> observations (which are NA) from the returned matrix.
seq_len	Data sequence length (i.e. number of past observations) used to predict each response quantile. By default, the training <code>fit_eqrn\$seq_len</code> is used.
device	(optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .

Value

Matrix of size `nrow(X)` times `prob_lvls_predict` (or `nrow(X)-seq_len` times `prob_lvls_predict` if `crop_predictions`) containing the conditional quantile estimates of the corresponding response observations at each probability level. Simplifies to a vector if `length(prob_lvls_predict)==1`.

EQRN_save

Save an EQRN object on disc

Description

Creates a folder named `name` and located in `path`, containing binary save files, so that the given "EQRN" object `fit_eqrn` can be loaded back in memory from disc using [EQRN_load\(\)](#).

Usage

```
EQRN_save(fit_eqrn, path, name = NULL, no_warning = TRUE)
```

Arguments

<code>fit_eqrn</code>	An "EQRN" object
<code>path</code>	Path to save folder as a string.
<code>name</code>	String name of the save.
<code>no_warning</code>	Whether to silence the warning raised if a save folder needed being created (bool).

Value

No return value.

excess_probability *Excess Probability Predictions*

Description

A generic function (method) for excess probability predictions from various fitted EQR models. The function invokes particular methods which depend on the class of the first argument.

Usage

```
excess_probability(object, ...)
```

Arguments

object	A model object for which excess probability prediction is desired.
...	additional model-specific arguments affecting the predictions produced. See the corresponding method documentation.

Value

The excess probability estimates from the given EQR model.

excess_probability.EQRN_iid
Tail excess probability prediction method using an EQRN_iid object

Description

Tail excess probability prediction method using an EQRN_iid object

Usage

```
## S3 method for class 'EQRN_iid'
excess_probability(object, ...)
```

Arguments

object	Fitted "EQRN_iid" object.
...	Arguments passed on to EQRN_excess_probability
val	Quantile value(s) used to estimate the conditional excess probability or cdf.
X	Matrix of covariates to predict the corresponding response's conditional excess probabilities.
intermediate_quantiles	Vector of intermediate conditional quantiles at level <code>fit_eqrn\$interm_lvl</code> .

`interm_lvl` Optional, checks that `interm_lvl == fit_eqrn$interm_lvl`.
`body_proba` Value to use when the predicted conditional probability is below `interm_lvl` (in which case it cannot be precisely assessed by the model). If "default" is given (the default), `paste0(">", 1-interm_lvl)` is used if `proba_type=="excess"`, and `paste0("<", interm_lvl)` is used if `proba_type=="cdf"`.
`proba_type` Whether to return the "excess" probability over `val` (default) or the "cdf" at `val`.
`device` (optional) A `torch::torch_device()`. Defaults to `default_device()`.

Details

See `EQRN_excess_probability()` for more details.

Value

Vector of probabilities (and possibly a few `body_proba` values if `val` is not large enough) of length `nrow(X)`.

excess_probability.EQRN_seq

Tail excess probability prediction method using an EQRN_iid object

Description

Tail excess probability prediction method using an `EQRN_iid` object

Usage

```
## S3 method for class 'EQRN_seq'
excess_probability(object, ...)
```

Arguments

<code>object</code>	Fitted "EQRN_seq" object.
<code>...</code>	Arguments passed on to <code>EQRN_excess_probability_seq</code>
<code>val</code>	Quantile value(s) used to estimate the conditional excess probability or cdf.
<code>X</code>	Matrix of covariates to predict the response's conditional excess probabilities.
<code>Y</code>	Response variable vector corresponding to the rows of <code>X</code> .
<code>intermediate_quantiles</code>	Vector of intermediate conditional quantiles at level <code>fit_eqrn\$interm_lvl</code> .
<code>interm_lvl</code>	Optional, checks that <code>interm_lvl == fit_eqrn\$interm_lvl</code> .
<code>crop_predictions</code>	Whether to crop out the first <code>seq_len</code> observations (which are NA) from the returned vector

`body_proba` Value to use when the predicted conditional probability is below `interm_lv1` (in which case it cannot be precisely assessed by the model). If "default" is given (the default), `paste0(">", 1-interm_lv1)` is used if `proba_type=="excess"`, and `paste0("<", interm_lv1)` is used if `proba_type=="cdf"`.

`proba_type` Whether to return the "excess" probability over `val` (default) or the "cdf" at `val`.

`seq_len` Data sequence length (i.e. number of past observations) used to predict each response quantile. By default, the training `fit_eqrn$seq_len` is used.

`device` (optional) A `torch::torch_device()`. Defaults to `default_device()`.

Details

See [EQRN_excess_probability_seq\(\)](#) for more details.

Value

Vector of probabilities (and possibly a few `body_proba` values if `val` is not large enough) of length `nrow(X)` (or `nrow(X)-seq_len` if `crop_predictions`).

FC_GPD_net

MLP module for GPD parameter prediction

Description

A fully-connected network (or multi-layer perception) as a `torch::nn_module`, designed for generalized Pareto distribution parameter prediction.

Usage

```
FC_GPD_net(
  D_in,
  Hidden_vect = c(5, 5, 5),
  activation = torch::nnf_sigmoid,
  p_drop = 0,
  shape_fixed = FALSE,
  device = EQRN::default_device()
)
```

Arguments

<code>D_in</code>	the input size (i.e. the number of features),
<code>Hidden_vect</code>	a vector of integers whose length determines the number of layers in the neural network and entries the number of neurons in each corresponding successive layer,
<code>activation</code>	the activation function for the hidden layers (should be either a callable function, preferably from the torch library),

<code>p_drop</code>	probability parameter for dropout before each hidden layer for regularization during training,
<code>shape_fixed</code>	whether the shape estimate depends on the covariates or not (bool),
<code>device</code>	a <code>torch::torch_device()</code> for an internal constant vector. Defaults to <code>default_device()</code> .

Details

The constructor allows specifying:

D_in the input size (i.e. the number of features),

Hidden_vect a vector of integers whose length determines the number of layers in the neural network and entries the number of neurons in each corresponding successive layer,

activation the activation function for the hidden layers (should be either a callable function, preferably from the torch library),

p_drop probability parameter for dropout before each hidden layer for regularization during training,

shape_fixed whether the shape estimate depends on the covariates or not (bool),

device a `torch::torch_device()` for an internal constant vector. Defaults to `default_device()`.

Value

The specified MLP GPD network as a `torch::nn_module`.

FC_GPD_SNN	<i>Self-normalized fully-connected network module for GPD parameter prediction</i>
------------	--

Description

A fully-connected self-normalizing network as a `torch::nn_module`, designed for generalized Pareto distribution parameter prediction.

Usage

```
FC_GPD_SNN(D_in, Hidden_vect = c(64, 64, 64), p_drop = 0.01)
```

Arguments

<code>D_in</code>	the input size (i.e. the number of features),
<code>Hidden_vect</code>	a vector of integers whose length determines the number of layers in the neural network and entries the number of neurons in each corresponding successive layer,
<code>p_drop</code>	probability parameter for the alpha-dropout before each hidden layer for regularization during training.

Details

The constructor allows specifying:

D_in the input size (i.e. the number of features),

Hidden_vect a vector of integers whose length determines the number of layers in the neural network and entries the number of neurons in each corresponding successive layer,

p_drop probability parameter for the alpha-dropout before each hidden layer for regularization during training.

Value

The specified SNN MLP GPD network as a `torch::nn_module`.

References

Gunter Klambauer, Thomas Unterthiner, Andreas Mayr, Sepp Hochreiter. Self-Normalizing Neural Networks. Advances in Neural Information Processing Systems 30 (NIPS 2017), 2017.

`fit_GPD_unconditional` *Maximum likelihood estimates for the GPD distribution using peaks over threshold*

Description

Maximum likelihood estimates for the GPD distribution using peaks over threshold

Usage

```
fit_GPD_unconditional(Y, interm_lvl = NULL, thresh_quantiles = NULL)
```

Arguments

<code>Y</code>	Vector of observations
<code>interm_lvl</code>	Probability level at which the empirical quantile should be used as the threshold, if <code>thresh_quantiles</code> is not given.
<code>thresh_quantiles</code>	Numerical value or numerical vector of the same length as <code>Y</code> representing either a fixed or a varying threshold, respectively.

Value

Named list containing:

<code>scale</code>	the GPD scale MLE,
<code>shape</code>	the GPD shape MLE,
<code>fit</code>	the fitted <code>ismev::gpd.fit()</code> object.

get_doFuture_operator *Get doFuture operator*

Description

Get doFuture operator

Usage

```
get_doFuture_operator(  
  strategy = c("sequential", "multisession", "multicore", "mixed")  
)
```

Arguments

strategy One of "sequential" (default), "multisession", "multicore", or "mixed".

Value

Returns the appropriate operator to use in a `foreach::foreach()` loop. The `%do%` operator is returned if `strategy=="sequential"`. Otherwise, the `%dopar%` operator is returned.

Examples

```
`%fun%` <- get_doFuture_operator("sequential")
```

get_excesses *Computes rescaled excesses over the conditional quantiles*

Description

Computes rescaled excesses over the conditional quantiles

Usage

```
get_excesses(  
  X = NULL,  
  y,  
  quantiles,  
  intermediate_q_feature = FALSE,  
  scale_features = FALSE,  
  X_scaling = NULL  
)
```

Arguments

<code>X</code>	A covariate matrix. Can be NULL if there are no covariates.
<code>y</code>	The response variable vector.
<code>quantiles</code>	The intermediate quantiles over which to compute the excesses of <code>y</code> .
<code>intermediate_q_feature</code>	Whether to use the intermediate quantiles as an additional covariate, by appending it to the <code>X</code> matrix (bool).
<code>scale_features</code>	Whether to rescale each input covariates to zero mean and unit variance before applying the network (recommended). If <code>X_scaling</code> is given, <code>X_scaling\$scaling</code> overrides <code>scale_features</code> .
<code>X_scaling</code>	Existing " <code>X_scaling</code> " object containing the precomputed mean and variance for each covariate. This enables reusing the scaling choice and parameters from the train set, if computing the excesses on a validation or test set, in order to avoid overfitting. This is performed automatically in the "EQRN" objects.

Value

Named list containing:

<code>Y_excesses</code>	the matrix of response excesses,
<code>X_excesses</code>	the (possibly rescaled and <code>q_feat</code> transformed) covariate matrix,
<code>X_scaling</code>	object of class " <code>X_scaling</code> " to use for consistent scaling on future datasets,
<code>excesses_ratio</code>	and the ratio of excesses for troubleshooting.

GPD_excess_probability

Tail excess probability prediction based on conditional GPD parameters

Description

Tail excess probability prediction based on conditional GPD parameters

Usage

```
GPD_excess_probability(
  val,
  sigma,
  xi,
  interm_threshold,
  threshold_p,
  body_proba = "default",
  proba_type = c("excess", "cdf")
)
```

Arguments

val	Quantile value(s) used to estimate the conditional excess probability or cdf.
sigma	Value(s) for the GPD scale parameter.
xi	Value(s) for the GPD shape parameter.
interm_threshold	Intermediate (conditional) quantile(s) at level threshold_p used as a (varying) threshold.
threshold_p	Probability level of the intermediate conditional quantiles interm_threshold.
body_proba	Value to use when the predicted conditional probability is below threshold_p (in which case it cannot be precisely assessed by the model). If "default" is given (the default), paste0(">", 1-threshold_p) is used if proba_type=="excess", and paste0("<", threshold_p) is used if proba_type=="cdf".
proba_type	Whether to return the "excess" probability over val (default) or the "cdf" at val.

Value

Vector of probabilities (and possibly a few body_proba values if val is not large enough) of the same length as the longest vector between val, sigma, xi and interm_threshold.

GPD_quantiles	<i>Compute extreme quantile from GPD parameters</i>
---------------	---

Description

Compute extreme quantile from GPD parameters

Usage

```
GPD_quantiles(p, p0, t_x0, sigma, xi)
```

Arguments

p	Probability level of the desired extreme quantile.
p0	Probability level of the (possibly varying) intermediate threshold/quantile.
t_x0	Value(s) of the (possibly varying) intermediate threshold/quantile.
sigma	Value(s) for the GPD scale parameter.
xi	Value(s) for the GPD shape parameter.

Value

The quantile value at probability level p.

install_backend	<i>Install Torch Backend Libraries</i>
-----------------	--

Description

This function can be called just after installing the EQRN package. Calling `EQRN::install_backend()` installs the necessary LibTorch and LibLantern backend libraries of the `torch` dependency by calling `torch::install_torch()`. See <https://torch.mlverse.org/docs/articles/installation> for more details and troubleshooting. Calling this function shouldn't be necessary in interactive environments, as loading EQRN (e.g. with `library(EQRN)` or with any `EQRN::fct()`) should prompt to do it automatically (via `.onLoad()`). This behaviour is inherited from the `torch` package.

Usage

```
install_backend(...)
```

Arguments

... Arguments passed to `torch::install_torch()`.

Value

No return value.

lagged_features	<i>Covariate lagged replication for temporal dependence</i>
-----------------	---

Description

Covariate lagged replication for temporal dependence

Usage

```
lagged_features(X, max_lag, drop_present = TRUE)
```

Arguments

`X` Covariate matrix.
`max_lag` Integer giving the maximum lag (i.e. the number of temporal dependence steps).
`drop_present` Whether to drop the "present" features (bool).

Value

Matrix with the original columns replicated, and shifted by `1:max_lag` if `drop_present==TRUE` (default) or by `0:max_lag` if `drop_present==FALSE`.

Examples

```
lagged_features(matrix(seq(20), ncol=2), max_lag=3, drop_present=TRUE)
```

last_elem	<i>Last element of a vector</i>
-----------	---------------------------------

Description

Returns the last element of the given vector in the most efficient way.

Usage

```
last_elem(x)
```

Arguments

x Vector.

Details

The last element is obtained using `x[length(x)]`, which is done in $O(1)$ and faster than, for example, any of `Rcpp::mylast(x)`, `tail(x, n=1)`, `dplyr::last(x)`, `x[end(x)[1]]`, and `rev(x)[1]`.

Value

The last element in the vector `x`.

Examples

```
last_elem(c(2, 6, 1, 4))
```

loss_GPD	<i>Generalized Pareto likelihood loss</i>
----------	---

Description

Generalized Pareto likelihood loss

Usage

```
loss_GPD(
  sigma,
  xi,
  y,
  rescaled = TRUE,
  interm_lvl = NULL,
  return_vector = FALSE
)
```

Arguments

sigma	Value(s) for the GPD scale parameter.
xi	Value(s) for the GPD shape parameter.
y	Vector of observations
rescaled	Whether y already is a vector of excesses (TRUE) or needs rescaling (FALSE).
interm_lvl	Probability level at which the empirical quantile should be used as the intermediate threshold to compute the excesses, if rescaled==FALSE.
return_vector	Whether to return the the vector of GPD losses for each observation instead of the negative log-likelihood (average loss).

Value

GPD negative log-likelihood of the GPD parameters over the sample of observations.

loss_GPD_tensor	<i>GPD tensor loss function for training a EQRN network</i>
-----------------	---

Description

GPD tensor loss function for training a EQRN network

Usage

```
loss_GPD_tensor(
  out,
  y,
  orthogonal_gpd = TRUE,
  shape_penalty = 0,
  prior_shape = NULL,
  return_agg = c("mean", "sum", "vector", "nanmean", "nansum")
)
```

Arguments

out	Batch tensor of GPD parameters output by the network.
y	Batch tensor of corresponding response variable.
orthogonal_gpd	Whether the network is supposed to regress in the orthogonal reparametrization of the GPD parameters (recommended).
shape_penalty	Penalty parameter for the shape estimate, to potentially regularize its variation from the fixed prior estimate.
prior_shape	Prior estimate for the shape, used only if shape_penalty>0.
return_agg	The return aggregation of the computed loss over the batch. Must be one of "mean", "sum", "vector", "nanmean", "nansum".

Value

The GPD loss over the batch between the network output and the observed responses as a `torch::Tensor`, whose dimensions depend on `return_agg`.

make_folds	<i>Create cross-validation folds</i>
------------	--------------------------------------

Description

Utility function to create folds of data, used in cross-validation procedures. The implementation is originally from the `gbx` R package

Usage

```
make_folds(y, num_folds, stratified = FALSE)
```

Arguments

<code>y</code>	Numerical vector of observations
<code>num_folds</code>	Number of folds to create.
<code>stratified</code>	Logical value. If TRUE, the folds are stratified along <code>rank(y)</code> .

Value

Vector of indices of the assigned folds for each observation.

Examples

```
make_folds(rnorm(30), 5)
```

mean_absolute_error	<i>Mean absolute error</i>
---------------------	----------------------------

Description

Mean absolute error

Usage

```
mean_absolute_error(
  y,
  y_hat,
  return_agg = c("mean", "sum", "vector"),
  na.rm = FALSE
)
```

Arguments

<code>y</code>	Vector of observations or ground-truths.
<code>y_hat</code>	Vector of predictions.
<code>return_agg</code>	Whether to return the "mean" (default), "sum", or "vector" of errors.
<code>na.rm</code>	A logical value indicating whether NA values should be stripped before the computation proceeds.

Value

The mean (or total or vectorial) absolute error between `y` and `y_hat`.

Examples

```
mean_absolute_error(c(2.3, 4.2, 1.8), c(2.2, 4.6, 1.7))
```

<code>mean_squared_error</code>	<i>Mean squared error</i>
---------------------------------	---------------------------

Description

Mean squared error

Usage

```
mean_squared_error(  
  y,  
  y_hat,  
  return_agg = c("mean", "sum", "vector"),  
  na.rm = FALSE  
)
```

Arguments

<code>y</code>	Vector of observations or ground-truths.
<code>y_hat</code>	Vector of predictions.
<code>return_agg</code>	Whether to return the "mean" (default), "sum", or "vector" of errors.
<code>na.rm</code>	A logical value indicating whether NA values should be stripped before the computation proceeds.

Value

The mean (or total or vectorial) squared error between `y` and `y_hat`.

Examples

```
mean_squared_error(c(2.3, 4.2, 1.8), c(2.2, 4.6, 1.7))
```

mts_dataset

*Dataset creator for sequential data***Description**

A `torch::dataset` object that can be initialized with sequential data, used to feed a recurrent network during training or prediction. It is used in `EQRN_fit_seq()` and corresponding predict functions, as well as in other recurrent methods such as `QRN_seq_fit()` and its predict functions. It can perform scaling of the response's past as a covariate, and compute excesses as a response when used in `EQRN_fit_seq()`. It also allows for fold separation or sequential discontinuity in the data.

Usage

```
mts_dataset(
  Y,
  X,
  seq_len,
  intermediate_quantiles = NULL,
  scale_Y = TRUE,
  fold_separation = NULL,
  sample_frac = 1,
  device = EQRN::default_device()
)
```

Arguments

<code>Y</code>	Response variable vector to model the extreme conditional quantile of, for training. Entries must be in sequential order.
<code>X</code>	Matrix of covariates, for training. Entries must be in sequential order.
<code>seq_len</code>	Data sequence length (i.e. number of past observations) used during training to predict each response quantile.
<code>intermediate_quantiles</code>	Vector of intermediate conditional quantiles at level <code>interm_lvl</code> .
<code>scale_Y</code>	Whether to rescale the response past, when considered as an input covariate, to zero mean and unit covariance before applying the network (recommended).
<code>fold_separation</code>	Fold separation index, when using concatenated folds as data.
<code>sample_frac</code>	Value between 0 and 1. If <code>sample_frac < 1</code> , a subsample of the data is used. Defaults to 1.
<code>device</code>	(optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .

Value

The `torch::dataset` containing the given data, to be used with a recurrent neural network.

```
multilevel_exceedance_proba_error
      Multilevel 'quantile_exceedance_proba_error'
```

Description

Multilevel version of [quantile_exceedance_proba_error\(\)](#).

Usage

```
multilevel_exceedance_proba_error(
  Probs,
  proba_levels = NULL,
  return_years = NULL,
  type_probs = c("cdf", "exceedance"),
  prefix = "",
  na.rm = FALSE,
  give_names = TRUE
)
```

Arguments

Probs	Matrix, whose columns give, for each proba_levels, the predicted probabilities to exceed or be smaller than a fixed quantile.
proba_levels	Vector of probability levels of the quantiles.
return_years	The probability levels can be given in term or return years instead. Only used if proba_levels is not given.
type_probs	Whether the predictions are the "cdf" (default) or "exceedance" probabilities.
prefix	A string prefix to add to the output's names (if give_names is TRUE).
na.rm	A logical value indicating whether NA values should be stripped before the computation proceeds.
give_names	Whether to name the output errors (bool).

Value

A vector of length `length(proba_levels)` giving the [quantile_exceedance_proba_error\(\)](#) calibration metric of each column of Probs at the corresponding proba_levels. If give_names is TRUE, the output vector is named `paste0(prefix, "exPrErr_q", proba_levels)` (or `paste0(prefix, "exPrErr_", return_years, "y")` if return_years are given instead of proba_levels).

multilevel_MAE	<i>Multilevel quantile MAEs</i>
----------------	---------------------------------

Description

Multilevel version of [mean_absolute_error\(\)](#).

Usage

```
multilevel_MAE(
  True_Q,
  Pred_Q,
  proba_levels,
  prefix = "",
  na.rm = FALSE,
  give_names = TRUE,
  sd = FALSE
)
```

Arguments

True_Q	Matrix of size n_obs times proba_levels, whose columns are the vectors of ground-truths at each proba_levels and each row corresponds to an observation or realisation.
Pred_Q	Matrix of the same size as True_Q, whose columns are the predictions at each proba_levels and each row corresponds to an observation or realisation.
proba_levels	Vector of probability levels at which the predictions were made. Must be of length ncol(Pred_Q).
prefix	A string prefix to add to the output's names (if give_names is TRUE).
na.rm	A logical value indicating whether NA values should be stripped before the computation proceeds.
give_names	Whether to name the output MAEs (bool).
sd	Whether to return the absolute error standard deviation (bool).

Value

A vector of length length(proba_levels) giving the mean absolute errors between each respective columns of True_Q and Pred_Q. If give_names is TRUE, the output vector is named paste0(prefix, "MAE_q", proba_levels). If sd==TRUE a named list is instead returned, containing the "MAEs" described above and "SDs", their standard deviations.

multilevel_MSE	<i>Multilevel quantile MSEs</i>
----------------	---------------------------------

Description

Multilevel version of `mean_squared_error()`.

Usage

```
multilevel_MSE(
  True_Q,
  Pred_Q,
  proba_levels,
  prefix = "",
  na.rm = FALSE,
  give_names = TRUE,
  sd = FALSE
)
```

Arguments

True_Q	Matrix of size n_obs times proba_levels, whose columns are the vectors of ground-truths at each proba_levels and each row corresponds to an observation or realisation.
Pred_Q	Matrix of the same size as True_Q, whose columns are the predictions at each proba_levels and each row corresponds to an observation or realisation.
proba_levels	Vector of probability levels at which the predictions were made. Must be of length ncol(Pred_Q).
prefix	A string prefix to add to the output's names (if give_names is TRUE).
na.rm	A logical value indicating whether NA values should be stripped before the computation proceeds.
give_names	Whether to name the output MSEs (bool).
sd	Whether to return the squared error standard deviation (bool).

Value

A vector of length `length(proba_levels)` giving the mean square errors between each respective columns of `True_Q` and `Pred_Q`. If `give_names` is TRUE, the output vector is named `paste0(prefix, "MSE_q", proba_levels)`. If `sd==TRUE` a named list is instead returned, containing the "MSEs" described above and "SDs", their standard deviations.

multilevel_pred_bias *Multilevel prediction bias*

Description

Multilevel version of `prediction_bias()`.

Usage

```
multilevel_pred_bias(  
  True_Q,  
  Pred_Q,  
  proba_levels,  
  square_bias = FALSE,  
  prefix = "",  
  na.rm = FALSE,  
  give_names = TRUE  
)
```

Arguments

True_Q	Matrix of size n_obs times proba_levels, whose columns are the vectors of ground-truths at each proba_levels and each row corresponds to an observation or realisation.
Pred_Q	Matrix of the same size as True_Q, whose columns are the predictions at each proba_levels and each row corresponds to an observation or realisation.
proba_levels	Vector of probability levels at which the predictions were made. Must be of length ncol(Pred_Q).
square_bias	Whether to return the square bias (bool); defaults to FALSE.
prefix	A string prefix to add to the output's names (if give_names is TRUE).
na.rm	A logical value indicating whether NA values should be stripped before the computation proceeds.
give_names	Whether to name the output MSEs (bool).

Value

A vector of length `length(proba_levels)` giving the (square) bias of each columns of predictions in `Pred_Q` for the respective `True_Q`. If `give_names` is TRUE, the output vector is named `paste0(prefix, "MSE_q", proba_levels)`.

multilevel_prop_below *Multilevel 'proportion_below'*

Description

Multilevel version of [proportion_below\(\)](#).

Usage

```
multilevel_prop_below(
  y,
  Pred_Q,
  proba_levels,
  prefix = "",
  na.rm = FALSE,
  give_names = TRUE
)
```

Arguments

<code>y</code>	Vector of observations.
<code>Pred_Q</code>	Matrix of of size <code>length(y)</code> times <code>proba_levels</code> , whose columns are the quantile predictions at each <code>proba_levels</code> and each row corresponds to an observation or realisation.
<code>proba_levels</code>	Vector of probability levels at which the predictions were made. Must be of length <code>ncol(Pred_Q)</code> .
<code>prefix</code>	A string prefix to add to the output's names (if <code>give_names</code> is <code>TRUE</code>).
<code>na.rm</code>	A logical value indicating whether NA values should be stripped before the computation proceeds.
<code>give_names</code>	Whether to name the output proportions (bool).

Value

A vector of length `length(proba_levels)` giving the proportion of observations below the predictions (`Pred_Q`) at each probability level. If `give_names` is `TRUE`, the output vector is named `paste0(prefix, "propBelow_q", proba_levels)`.

multilevel_q_loss	<i>Multilevel quantile losses</i>
-------------------	-----------------------------------

Description

Multilevel version of [quantile_loss\(\)](#).

Usage

```
multilevel_q_loss(
  y,
  Pred_Q,
  proba_levels,
  prefix = "",
  na.rm = FALSE,
  give_names = TRUE
)
```

Arguments

<code>y</code>	Vector of observations.
<code>Pred_Q</code>	Matrix of of size <code>length(y)</code> times <code>proba_levels</code> , whose columns are the quantile predictions at each <code>proba_levels</code> and each row corresponds to an observation or realisation.
<code>proba_levels</code>	Vector of probability levels at which the predictions were made. Must be of length <code>ncol(Pred_Q)</code> .
<code>prefix</code>	A string prefix to add to the output's names (if <code>give_names</code> is TRUE).
<code>na.rm</code>	A logical value indicating whether NA values should be stripped before the computation proceeds.
<code>give_names</code>	Whether to name the output quantile errors (bool).

Value

A vector of length `length(proba_levels)` giving the average quantile losses between each column of `Pred_Q` and the observations. If `give_names` is TRUE, the output vector is named `paste0(prefix, "qloss_q", proba_levels)`.

```
multilevel_q_pred_error
      Multilevel 'quantile_prediction_error'
```

Description

Multilevel version of [quantile_prediction_error\(\)](#).

Usage

```
multilevel_q_pred_error(
  y,
  Pred_Q,
  proba_levels,
  prefix = "",
  na.rm = FALSE,
  give_names = TRUE
)
```

Arguments

<code>y</code>	Vector of observations.
<code>Pred_Q</code>	Matrix of of size <code>length(y)</code> times <code>proba_levels</code> , whose columns are the quantile predictions at each <code>proba_levels</code> and each row corresponds to an observation or realisation.
<code>proba_levels</code>	Vector of probability levels at which the predictions were made. Must be of length <code>ncol(Pred_Q)</code> .
<code>prefix</code>	A string prefix to add to the output's names (if <code>give_names</code> is TRUE).
<code>na.rm</code>	A logical value indicating whether NA values should be stripped before the computation proceeds.
<code>give_names</code>	Whether to name the output errors (bool).

Value

A vector of length `length(proba_levels)` giving the quantile prediction error calibration metrics between each column of `Pred_Q` and the observations. If `give_names` is TRUE, the output vector is named `paste0(prefix, "qPredErr_q", proba_levels)`.

multilevel_resid_var *Multilevel residual variance*

Description

Multilevel version of [prediction_residual_variance\(\)](#).

Usage

```
multilevel_resid_var(  
  True_Q,  
  Pred_Q,  
  proba_levels,  
  prefix = "",  
  na.rm = FALSE,  
  give_names = TRUE  
)
```

Arguments

True_Q	Matrix of size n_obs times proba_levels, whose columns are the vectors of ground-truths at each proba_levels and each row corresponds to an observation or realisation.
Pred_Q	Matrix of the same size as True_Q, whose columns are the predictions at each proba_levels and each row corresponds to an observation or realisation.
proba_levels	Vector of probability levels at which the predictions were made. Must be of length ncol(Pred_Q).
prefix	A string prefix to add to the output's names (if give_names is TRUE).
na.rm	A logical value indicating whether NA values should be stripped before the computation proceeds.
give_names	Whether to name the output MSEs (bool).

Value

A vector of length length(proba_levels) giving the residual variances of each columns of predictions in Pred_Q for the respective True_Q. If give_names is TRUE, the output vector is named paste0(prefix, "MSE_q", proba_levels).

multilevel_R_squared *Multilevel R squared*

Description

Multilevel version of `R_squared()`.

Usage

```
multilevel_R_squared(
  True_Q,
  Pred_Q,
  proba_levels,
  prefix = "",
  na.rm = FALSE,
  give_names = TRUE
)
```

Arguments

<code>True_Q</code>	Matrix of size <code>n_obs</code> times <code>proba_levels</code> , whose columns are the vectors of ground-truths at each <code>proba_levels</code> and each row corresponds to an observation or realisation.
<code>Pred_Q</code>	Matrix of the same size as <code>True_Q</code> , whose columns are the predictions at each <code>proba_levels</code> and each row corresponds to an observation or realisation.
<code>proba_levels</code>	Vector of probability levels at which the predictions were made. Must be of length <code>ncol(Pred_Q)</code> .
<code>prefix</code>	A string prefix to add to the output's names (if <code>give_names</code> is <code>TRUE</code>).
<code>na.rm</code>	A logical value indicating whether NA values should be stripped before the computation proceeds.
<code>give_names</code>	Whether to name the output MSEs (bool).

Value

A vector of length `length(proba_levels)` giving the R squared coefficient of determination of each columns of predictions in `Pred_Q` for the respective `True_Q`. If `give_names` is `TRUE`, the output vector is named `paste0(prefix, "MSE_q", proba_levels)`.

perform_scaling	<i>Performs feature scaling without overfitting</i>
-----------------	---

Description

Performs feature scaling without overfitting

Usage

```
perform_scaling(X, X_scaling = NULL, scale_features = TRUE, stat_attr = FALSE)
```

Arguments

X	A covariate matrix.
X_scaling	Existing "X_scaling" object containing the precomputed mean and variance for each covariate. This enables reusing the scaling choice and parameters from the train set, if computing the excesses on a validation or test set, in order to avoid overfitting. This is performed automatically in the "EQRN" objects.
scale_features	Whether to rescale each input covariates to zero mean and unit variance before applying the model (recommended). If X_scaling is given, X_scaling\$scaling overrides scale_features.
stat_attr	DEPRECATED. Whether to keep attributes in the returned covariate matrix itself.

Value

Named list containing:

X_excesses	the (possibly rescaled and q_feat transformed) covariate matrix,
X_scaling	object of class "X_scaling" to use for consistent scaling on future datasets.

predict.EQRN_iid	<i>Predict method for an EQRN_iid fitted object</i>
------------------	---

Description

Predict method for an EQRN_iid fitted object

Usage

```
## S3 method for class 'EQRN_iid'
predict(object, ...)
```

Arguments

object Fitted "EQRN_iid" object.
 ... Arguments passed on to [EQRN_predict](#)
 X Matrix of covariates to predict the corresponding response's conditional quantiles.
 prob_lvls_predict Vector of probability levels at which to predict the conditional quantiles.
 intermediate_quantiles Vector of intermediate conditional quantiles at level `fit_eqrn$interm_lv1`.
 interm_lv1 Optional, checks that `interm_lv1 == fit_eqrn$interm_lv1`.
 device (optional) A `torch::torch_device()`. Defaults to `default_device()`.

Details

See [EQRN_predict\(\)](#) for more details.

Value

Matrix of size `nrow(X)` times `prob_lvls_predict` containing the conditional quantile estimates of the response associated to each covariate observation at each probability level. Simplifies to a vector if `length(prob_lvls_predict)==1`.

`predict.EQRN_seq` *Predict method for an EQRN_seq fitted object*

Description

Predict method for an EQRN_seq fitted object

Usage

```
## S3 method for class 'EQRN_seq'
predict(object, ...)
```

Arguments

object Fitted "EQRN_seq" object.
 ... Arguments passed on to [EQRN_predict_seq](#)
 X Matrix of covariates to predict the corresponding response's conditional quantiles.
 Y Response variable vector corresponding to the rows of X.
 prob_lvls_predict Vector of probability levels at which to predict the conditional quantiles.
 intermediate_quantiles Vector of intermediate conditional quantiles at level `fit_eqrn$interm_lv1`.

interm_lvl Optional, checks that `interm_lvl == fit_eqrn$interm_lvl`.

crop_predictions Whether to crop out the first `seq_len` observations (which are NA) from the returned matrix.

seq_len Data sequence length (i.e. number of past observations) used to predict each response quantile. By default, the training `fit_eqrn$seq_len` is used.

device (optional) A `torch::torch_device()`. Defaults to `default_device()`.

Details

See [EQRN_predict_seq\(\)](#) for more details.

Value

Matrix of size `nrow(X)` times `prob_lvls_predict` (or `nrow(X)-seq_len` times `prob_lvls_predict` if `crop_predictions`) containing the conditional quantile estimates of the corresponding response observations at each probability level. Simplifies to a vector if `length(prob_lvls_predict)==1`.

predict.QRN_seq	<i>Predict method for a QRN_seq fitted object</i>
-----------------	---

Description

Predict method for a QRN_seq fitted object

Usage

```
## S3 method for class 'QRN_seq'
predict(object, ...)
```

Arguments

object	Fitted "QRN_seq" object.
...	Arguments passed on to QRN_seq_predict
	X Matrix of covariates to predict the corresponding response's conditional quantiles.
	Y Response variable vector corresponding to the rows of X.
	q_level Optional, checks that <code>q_level == fit_qrn_ts\$interm_lvl</code> .
	crop_predictions Whether to crop out the first <code>seq_len</code> observations (which are NA) from the returned matrix.
	device (optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .

Details

See [QRN_seq_predict\(\)](#) for more details.

Value

Matrix of size $nrow(X)$ times 1 (or $nrow(X)-seq_len$ times 1 if `crop_predictions`) containing the conditional quantile estimates of the corresponding response observations.

prediction_bias	<i>Prediction bias</i>
-----------------	------------------------

Description

Prediction bias

Usage

```
prediction_bias(y, y_hat, square_bias = FALSE, na.rm = FALSE)
```

Arguments

y	Vector of observations or ground-truths.
y_hat	Vector of predictions.
square_bias	Whether to return the square bias (bool); defaults to FALSE.
na.rm	A logical value indicating whether NA values should be stripped before the computation proceeds.

Value

The (square) bias of the predictions `y_hat` for `y`.

Examples

```
prediction_bias(c(2.3, 4.2, 1.8), c(2.2, 4.6, 1.7))
```

prediction_residual_variance	<i>Prediction residual variance</i>
------------------------------	-------------------------------------

Description

Prediction residual variance

Usage

```
prediction_residual_variance(y, y_hat, na.rm = FALSE)
```

Arguments

y	Vector of observations or ground-truths.
y_hat	Vector of predictions.
na.rm	A logical value indicating whether NA values should be stripped before the computation proceeds.

Value

The residual variance of the predictions y_hat for y.

Examples

```
prediction_residual_variance(c(2.3, 4.2, 1.8), c(2.2, 4.6, 1.7))
```

```
predict_GPD_semiconditional
```

Predict semi-conditional extreme quantiles using peaks over threshold

Description

Predict semi-conditional extreme quantiles using peaks over threshold

Usage

```
predict_GPD_semiconditional(
  Y,
  interm_lvl,
  thresh_quantiles,
  interm_quantiles_test = thresh_quantiles,
  prob_lvls_predict = c(0.99)
)
```

Arguments

Y	Vector of ("training") observations.
interm_lvl	Probability level at which the empirical quantile should be used as the intermediate threshold.
thresh_quantiles	Numerical vector of the same length as Y representing the varying intermediate threshold on the train set.
interm_quantiles_test	Numerical vector of the same length as Y representing the varying intermediate threshold used for prediction on the test set.
prob_lvls_predict	Probability levels at which to predict the extreme semi-conditional quantiles.

Value

Named list containing:

predictions	matrix of dimension <code>length(interm_quantiles_test)</code> times <code>length(prob_lvls_predict)</code> containing the estimated extreme quantile at levels <code>quantile</code> , for each <code>interm_quantiles_test</code> ,
pars	matrix of dimension <code>n_test</code> times 2 containing the two GPD parameter MLEs, repeated <code>length(interm_quantiles_test)</code> times.

predict_unconditional_quantiles

Predict unconditional extreme quantiles using peaks over threshold

Description

Predict unconditional extreme quantiles using peaks over threshold

Usage

```
predict_unconditional_quantiles(interm_lvl, quantiles = c(0.99), Y, ntest = 1)
```

Arguments

interm_lvl	Probability level at which the empirical quantile should be used as the intermediate threshold.
quantiles	Probability levels at which to predict the extreme quantiles.
Y	Vector of ("training") observations.
ntest	Number of "test" observations.

Value

Named list containing:

predictions	matrix of dimension <code>ntest</code> times <code>length(quantiles)</code> containing the estimated extreme quantile at levels <code>quantile</code> , repeated <code>ntest</code> times,
pars	matrix of dimension <code>ntest</code> times 2 containing the two GPD parameter MLEs, repeated <code>ntest</code> times.
threshold	The threshold for the peaks-over-threshold GPD model. It is the empirical quantile of <code>Y</code> at level <code>interm_lvl</code> , i.e. <code>stats::quantile(Y, interm_lvl)</code> .

process_features *Feature processor for EQRN*

Description

Feature processor for EQRN

Usage

```
process_features(
  X,
  intermediate_q_feature,
  intermediate_quantiles = NULL,
  X_scaling = NULL,
  scale_features = TRUE
)
```

Arguments

X A covariate matrix.

intermediate_q_feature Whether to use the intermediate quantiles as an additional covariate, by appending it to the X matrix (bool).

intermediate_quantiles The intermediate conditional quantiles.

X_scaling Existing "X_scaling" object containing the precomputed mean and variance for each covariate. This enables reusing the scaling choice and parameters from the train set, if computing the excesses on a validation or test set, in order to avoid overfitting. This is performed automatically in the "EQRN" objects.

scale_features Whether to rescale each input covariates to zero mean and unit variance before applying the network (recommended). If X_scaling is given, X_scaling\$scaling overrides scale_features.

Value

Named list containing:

X_excesses the (possibly rescaled and q_feat transformed) covariate matrix,

X_scaling object of class "X_scaling" to use for consistent scaling on future datasets.

proportion_below	<i>Proportion of observations below conditional quantile vector</i>
------------------	---

Description

Proportion of observations below conditional quantile vector

Usage

```
proportion_below(y, Q_hat, na.rm = FALSE)
```

Arguments

y	Vector of observations.
Q_hat	Vector of predicted quantiles.
na.rm	A logical value indicating whether NA values should be stripped before the computation proceeds.

Value

The proportion of observation below the predictions.

Examples

```
proportion_below(c(2.3, 4.2, 1.8), c(2.9, 5.6, 1.7))
```

QRNN_RNN_net	<i>Recurrent quantile regression neural network module</i>
--------------	--

Description

A recurrent neural network as a `torch::nn_module`, designed for quantile regression.

Usage

```
QRNN_RNN_net(  
  type = c("lstm", "gru"),  
  nb_input_features,  
  hidden_size,  
  num_layers = 1,  
  dropout = 0  
)
```

Arguments

type	the type of recurrent architecture, can be one of "lstm" (default) or "gru",
nb_input_features	the input size (i.e. the number of features),
hidden_size	the dimension of the hidden latent state variables in the recurrent network,
num_layers	the number of recurrent layers,
dropout	probability parameter for dropout before each hidden layer for regularization during training.

Details

The constructor allows specifying:

type the type of recurrent architecture, can be one of "lstm" (default) or "gru",

nb_input_features the input size (i.e. the number of features),

hidden_size the dimension of the hidden latent state variables in the recurrent network,

num_layers the number of recurrent layers,

dropout probability parameter for dropout before each hidden layer for regularization during training.

Value

The specified recurrent QRN as a `torch::nn::module`.

QRN_fit_multiple	<i>Wrapper for fitting a recurrent QRN with restart for stability</i>
------------------	---

Description

Wrapper for fitting a recurrent QRN with restart for stability

Usage

```
QRN_fit_multiple(
  X,
  y,
  q_level,
  number_fits = 3,
  ...,
  seed = NULL,
  data_type = c("seq", "iid")
)
```

Arguments

X	Matrix of covariates, for training.
y	Response variable vector to model the conditional quantile of, for training.
q_level	Probability level of the desired conditional quantiles to predict.
number_fits	Number of restarts.
...	Other parameters given to <code>QRN_seq_fit()</code> .
seed	Integer random seed for reproducibility in network weight initialization.
data_type	Type of data dependence, must be one of "iid" (for iid observations) or "seq" (for sequentially dependent observations). For the moment, only "seq" is accepted.

Value

An QRN object of classes `c("QRN_seq", "QRN")`, containing the fitted network, as well as all the relevant information for its usage in other functions.

QRN_seq_fit	<i>Recurrent QRN fitting function</i>
-------------	---------------------------------------

Description

Used to fit a recurrent quantile regression neural network on a data sample. Use the `QRN_fit_multiple()` wrapper instead, with `data_type="seq"`, for better stability using fitting restart.

Usage

```
QRN_seq_fit(
  X,
  Y,
  q_level,
  hidden_size = 10,
  num_layers = 1,
  rnn_type = c("lstm", "gru"),
  p_drop = 0,
  learning_rate = 1e-04,
  L2_pen = 0,
  seq_len = 10,
  scale_features = TRUE,
  n_epochs = 10000,
  batch_size = 256,
  X_valid = NULL,
  Y_valid = NULL,
  lr_decay = 1,
  patience_decay = n_epochs,
  min_lr = 0,
```

```

    patience_stop = n_epochs,
    tol = 1e-04,
    fold_separation = NULL,
    warm_start_path = NULL,
    patience_lag = 5,
    optim_met = "adam",
    seed = NULL,
    verbose = 2,
    device = default_device()
)

```

Arguments

X	Matrix of covariates, for training. Entries must be in sequential order.
Y	Response variable vector to model the conditional quantile of, for training. Entries must be in sequential order.
q_level	Probability level of the desired conditional quantiles to predict.
hidden_size	Dimension of the hidden latent state variables in the recurrent network.
num_layers	Number of recurrent layers.
rnn_type	Type of recurrent architecture, can be one of "lstm" (default) or "gru".
p_drop	Probability parameter for dropout before each hidden layer for regularization during training.
learning_rate	Initial learning rate for the optimizer during training of the neural network.
L2_pen	L2 weight penalty parameter for regularization during training.
seq_len	Data sequence length (i.e. number of past observations) used during training to predict each response quantile.
scale_features	Whether to rescale each input covariates to zero mean and unit covariance before applying the network (recommended).
n_epochs	Number of training epochs.
batch_size	Batch size used during training.
X_valid	Covariates in a validation set, or NULL. Entries must be in sequential order. Used for monitoring validation loss during training, enabling learning-rate decay and early stopping.
Y_valid	Response variable in a validation set, or NULL. Entries must be in sequential order. Used for monitoring validation loss during training, enabling learning-rate decay and early stopping.
lr_decay	Learning rate decay factor.
patience_decay	Number of epochs of non-improving validation loss before a learning-rate decay is performed.
min_lr	Minimum learning rate, under which no more decay is performed.
patience_stop	Number of epochs of non-improving validation loss before early stopping is performed.

tol	Tolerance for stopping training, in case of no significant training loss improvements.
fold_separation	Index of fold separation or sequential discontinuity in the data.
warm_start_path	Path of a saved network using <code>torch::torch_save()</code> , to load back for a warm start.
patience_lag	The validation loss is considered to be non-improving if it is larger than on any of the previous <code>patience_lag</code> epochs.
optim_met	DEPRECATED. Optimization algorithm to use during training. "adam" is the default.
seed	Integer random seed for reproducibility in network weight initialization.
verbose	Amount of information printed during training (0:nothing, 1:most important, 2:everything).
device	(optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .

Value

An QRN object of classes `c("QRN_seq", "QRN")`, containing the fitted network, as well as all the relevant information for its usage in other functions.

QRN_seq_predict	<i>Predict function for a QRN_seq fitted object</i>
-----------------	---

Description

Predict function for a QRN_seq fitted object

Usage

```
QRN_seq_predict(
  fit_qrn_ts,
  X,
  Y,
  q_level = fit_qrn_ts$interm_lvl,
  crop_predictions = FALSE,
  device = default_device()
)
```

Arguments

fit_qrn_ts	Fitted "QRN_seq" object.
X	Matrix of covariates to predict the corresponding response's conditional quantiles.
Y	Response variable vector corresponding to the rows of X.

q_level	Optional, checks that <code>q_level == fit_qrn_ts\$interm_lvl</code> .
crop_predictions	Whether to crop out the first <code>seq_len</code> observations (which are NA) from the returned matrix.
device	(optional) A <code>torch::torch_device()</code> . Defaults to <code>default_device()</code> .

Value

Matrix of size `nrow(X)` times 1 (or `nrow(X)-seq_len` times 1 if `crop_predictions`) containing the conditional quantile estimates of the corresponding response observations.

QRN_seq_predict_foldwise

Foldwise fit-predict function using a recurrent QRN

Description

Foldwise fit-predict function using a recurrent QRN

Usage

```
QRN_seq_predict_foldwise(  
  X,  
  y,  
  q_level,  
  n_folds = 3,  
  number_fits = 3,  
  seq_len = 10,  
  seed = NULL,  
  ...  
)
```

Arguments

X	Matrix of covariates, for training. Entries must be in sequential order.
y	Response variable vector to model the conditional quantile of, for training. Entries must be in sequential order.
q_level	Probability level of the desired conditional quantiles to predict.
n_folds	Number of folds.
number_fits	Number of restarts, for stability.
seq_len	Data sequence length (i.e. number of past observations) used during training to predict each response quantile.
seed	Integer random seed for reproducibility in network weight initialization.
...	Other parameters given to <code>QRN_seq_fit()</code> .

Value

A named list containing the foldwise predictions and fits. It namely contains:

predictions	the numerical vector of quantile predictions for each observation entry in y,
fits	a list containing the "QRN_seq" fitted networks for each fold,
cuts	the fold cuts indices,
folds	a list of lists containing the train indices, validation indices and fold separations as a list for each fold setup,
n_folds	number of folds,
q_level	probability level of the predicted quantiles,
train_losses	the vector of train losses on each fold,
valid_losses	the vector of validation losses on each fold,
min_valid_losses	the minimal validation losses obtained on each fold,
min_valid_e	the epoch index of the minimal validation losses obtained on each fold.

QRN_seq_predict_foldwise_sep

Single-fold foldwise fit-predict function using a recurrent QRN

Description

Separated single-fold version of [QRN_seq_predict_foldwise\(\)](#), for computation purposes.

Usage

```
QRN_seq_predict_foldwise_sep(
  X,
  y,
  q_level,
  n_folds = 3,
  fold_todo = 1,
  number_fits = 3,
  seq_len = 10,
  seed = NULL,
  ...
)
```

Arguments

X	Matrix of covariates, for training. Entries must be in sequential order.
y	Response variable vector to model the conditional quantile of, for training. Entries must be in sequential order.
q_level	Probability level of the desired conditional quantiles to predict.

n_folds	Number of folds.
fold_todo	Index of the fold to do (integer in 1:n_folds).
number_fits	Number of restarts, for stability.
seq_len	Data sequence length (i.e. number of past observations) used during training to predict each response quantile.
seed	Integer random seed for reproducibility in network weight initialization.
...	Other parameters given to <code>QRN_seq_fit()</code> .

Value

A named list containing the foldwise predictions and fits. It namely contains:

predictions	the numerical vector of quantile predictions for each observation entry in <code>y</code> ,
fits	a list containing the "QRN_seq" fitted networks for each fold,
cuts	the fold cuts indices,
folds	a list of lists containing the train indices, validation indices and fold separations as a list for each fold setup,
n_folds	number of folds,
q_level	probability level of the predicted quantiles,
train_losses	the vector of train losses on each fold,
valid_losses	the vector of validation losses on each fold,
min_valid_losses	the minimal validation losses obtained on each fold,
min_valid_e	the epoch index of the minimal validation losses obtained on each fold.

quantile_exceedance_proba_error

Quantile exceedance probability prediction calibration error

Description

Quantile exceedance probability prediction calibration error

Usage

```
quantile_exceedance_proba_error(
  Probs,
  prob_level = NULL,
  return_years = NULL,
  type_probs = c("cdf", "exceedance"),
  na.rm = FALSE
)
```

Arguments

Probs	Predicted probabilities to exceed or be smaller than a fixed quantile.
prob_level	Probability level of the quantile.
return_years	The probability level can be given in term or return years instead. Only used if prob_level is not given.
type_probs	Whether the predictions are the "cdf" (default) or "exceedance" probabilities.
na.rm	A logical value indicating whether NA values should be stripped before the computation proceeds.

Value

The calibration metric for the predicted probabilities.

Examples

```
quantile_exceedance_proba_error(c(0.1, 0.3, 0.2), prob_level=0.8)
```

quantile_loss	<i>Quantile loss</i>
---------------	----------------------

Description

Quantile loss

Usage

```
quantile_loss(
  y,
  y_hat,
  q,
  return_agg = c("mean", "sum", "vector"),
  na.rm = FALSE
)
```

Arguments

y	Vector of observations.
y_hat	Vector of predicted quantiles at probability level q.
q	Probability level of the predicted quantile.
return_agg	Whether to return the "mean" (default), "sum", or "vector" of losses.
na.rm	A logical value indicating whether NA values should be stripped before the computation proceeds.

Value

The mean (or total or vectorial) quantile loss between y and y_hat at level q.

Examples

```
quantile_loss(c(2.3, 4.2, 1.8), c(2.9, 5.6, 2.7), q=0.8)
```

quantile_loss_tensor *Tensor quantile loss function for training a QRN network*

Description

Tensor quantile loss function for training a QRN network

Usage

```
quantile_loss_tensor(
  out,
  y,
  q = 0.5,
  return_agg = c("mean", "sum", "vector", "nanmean", "nansum")
)
```

Arguments

out	Batch tensor of the quantile output by the network.
y	Batch tensor of corresponding response variable.
q	Probability level of the predicted quantile
return_agg	The return aggregation of the computed loss over the batch. Must be one of "mean", "sum", "vector", "nanmean", "nansum".

Value

The quantile loss over the batch between the network output and the observed responses as a torch::Tensor, whose dimensions depend on return_agg.

quantile_prediction_error
Quantile prediction calibration error

Description

Quantile prediction calibration error

Usage

```
quantile_prediction_error(y, Q_hat, prob_level, na.rm = FALSE)
```

Arguments

<code>y</code>	Vector of observations.
<code>Q_hat</code>	Vector of predicted quantiles at probability level <code>prob_level</code> .
<code>prob_level</code>	Probability level of the predicted quantile.
<code>na.rm</code>	A logical value indicating whether NA values should be stripped before the computation proceeds.

Value

The quantile prediction error calibration metric.

Examples

```
quantile_prediction_error(c(2.3, 4.2, 1.8), c(2.9, 5.6, 2.7), prob_level=0.8)
```

Recurrent_GPD_net *Recurrent network module for GPD parameter prediction*

Description

A recurrent neural network as a `torch::nn_module`, designed for generalized Pareto distribution parameter prediction, with sequential dependence.

Usage

```
Recurrent_GPD_net(
  type = c("lstm", "gru"),
  nb_input_features,
  hidden_size,
  num_layers = 1,
  dropout = 0,
  shape_fixed = FALSE,
  device = EQRN::default_device()
)
```

Arguments

<code>type</code>	the type of recurrent architecture, can be one of "lstm" (default) or "gru",
<code>nb_input_features</code>	the input size (i.e. the number of features),
<code>hidden_size</code>	the dimension of the hidden latent state variables in the recurrent network,
<code>num_layers</code>	the number of recurrent layers,
<code>dropout</code>	probability parameter for dropout before each hidden layer for regularization during training,
<code>shape_fixed</code>	whether the shape estimate depends on the covariates or not (bool),
<code>device</code>	a <code>torch::torch_device()</code> for an internal constant vector. Defaults to <code>default_device()</code> .

Details

The constructor allows specifying:

type the type of recurrent architecture, can be one of "lstm" (default) or "gru",

nb_input_features the input size (i.e. the number of features),

hidden_size the dimension of the hidden latent state variables in the recurrent network,

num_layers the number of recurrent layers,

dropout probability parameter for dropout before each hidden layer for regularization during training,

shape_fixed whether the shape estimate depends on the covariates or not (bool),

device a `torch::torch_device()` for an internal constant vector. Defaults to `default_device()`.

Value

The specified recurrent GPD network as a `torch::nn_module`.

roundm

Mathematical number rounding

Description

This function rounds numbers in the mathematical sense, as opposed to the base R function `round()` that rounds 'to the even digit'.

Usage

```
roundm(x, decimals = 0)
```

Arguments

`x` Vector of numerical values to round.

`decimals` Integer indicating the number of decimal places to be used.

Value

A vector containing the entries of `x`, rounded to `decimals` decimals.

Examples

```
roundm(2.25, 1)
```

R_squared	<i>R squared</i>
-----------	------------------

Description

The coefficient of determination, often called R squared, is the proportion of data variance explained by the predictions.

Usage

```
R_squared(y, y_hat, na.rm = FALSE)
```

Arguments

y	Vector of observations or ground-truths.
y_hat	Vector of predictions.
na.rm	A logical value indicating whether NA values should be stripped before the computation proceeds.

Value

The R squared of the predictions y_hat for y.

Examples

```
R_squared(c(2.3, 4.2, 1.8), c(2.2, 4.6, 1.7))
```

safe_save_rds	<i>Safe RDS save</i>
---------------	----------------------

Description

Safe version of [saveRDS\(\)](#). If the given save path (i.e. `dirname(file_path)`) does not exist, it is created instead of raising an error.

Usage

```
safe_save_rds(object, file_path, recursive = TRUE, no_warning = FALSE)
```

Arguments

object	R variable or object to save on disk.
file_path	Path and name of the save file, as a string.
recursive	Should elements of the path other than the last be created? If TRUE, behaves like the Unix command <code>mkdir -p</code> .
no_warning	Whether to cancel the warning issued if a directory is created (bool).

Value

No return value.

Examples

```
safe_save_rds(c(1, 2, 8), "./some_folder/my_new_folder/my_vector.rds")
```

```
semiconditional_train_valid_GPD_loss
```

Semi-conditional GPD MLEs and their train-validation likelihoods

Description

Semi-conditional GPD MLEs and their train-validation likelihoods

Usage

```
semiconditional_train_valid_GPD_loss(  
  Y_train,  
  Y_valid,  
  interm_quant_train,  
  interm_quant_valid  
)
```

Arguments

<code>Y_train</code>	Vector of "training" observations on which to estimate the MLEs.
<code>Y_valid</code>	Vector of "validation" observations, on which to estimate the out of training sample GPD loss.
<code>interm_quant_train</code>	Vector of intermediate quantiles serving as a varying threshold for each training observation.
<code>interm_quant_valid</code>	Vector of intermediate quantiles serving as a varying threshold for each validation observation.

Value

Named list containing:

<code>scale</code>	GPD scale MLE inferred from the train set,
<code>shape</code>	GPD shape MLE inferred from the train set,
<code>train_loss</code>	the negative log-likelihoods of the MLEs over the training samples,
<code>valid_loss</code>	the negative log-likelihoods of the MLEs over the validation samples.

Separated_GPD_SNN	<i>Self-normalized separated network module for GPD parameter prediction</i>
-------------------	--

Description

A parameter-separated self-normalizing network as a `torch::nn_module`, designed for generalized Pareto distribution parameter prediction.

Usage

```
Separated_GPD_SNN(
  D_in,
  Hidden_vect_scale = c(64, 64, 64),
  Hidden_vect_shape = c(5, 3),
  p_drop = 0.01
)
```

Arguments

<code>D_in</code>	the input size (i.e. the number of features),
<code>Hidden_vect_scale</code>	a vector of integers whose length determines the number of layers in the sub-network for the scale parameter and entries the number of neurons in each corresponding successive layer,
<code>Hidden_vect_shape</code>	a vector of integers whose length determines the number of layers in the sub-network for the shape parameter and entries the number of neurons in each corresponding successive layer,
<code>p_drop</code>	probability parameter for the alpha-dropout before each hidden layer for regularization during training.

Details

The constructor allows specifying:

D_in the input size (i.e. the number of features),

Hidden_vect_scale a vector of integers whose length determines the number of layers in the sub-network for the scale parameter and entries the number of neurons in each corresponding successive layer,

Hidden_vect_shape a vector of integers whose length determines the number of layers in the sub-network for the shape parameter and entries the number of neurons in each corresponding successive layer,

p_drop probability parameter for the alpha-dropout before each hidden layer for regularization during training.

Value

The specified parameter-separated SNN MLP GPD network as a `torch::nn_module`.

References

Gunter Klambauer, Thomas Unterthiner, Andreas Mayr, Sepp Hochreiter. Self-Normalizing Neural Networks. Advances in Neural Information Processing Systems 30 (NIPS 2017), 2017.

set_doFuture_strategy *Set a doFuture execution strategy*

Description

Set a doFuture execution strategy

Usage

```
set_doFuture_strategy(
  strategy = c("sequential", "multisession", "multicore", "mixed"),
  n_workers = NULL
)
```

Arguments

strategy	One of "sequential" (default), "multisession", "multicore", or "mixed".
n_workers	A positive numeric scalar or a function specifying the maximum number of parallel futures that can be active at the same time before blocking. If a function, it is called without arguments when the future is created and its value is used to configure the workers. The function should return a numeric scalar. Defaults to <code>future::availableCores()-1</code> if NULL (default), with "multicore" constraint in the relevant case. Ignored if <code>strategy=="sequential"</code> .

Value

The appropriate `get_doFuture_operator()` operator to use in a `foreach::foreach()` loop. The `%do%` operator is returned if `strategy=="sequential"`. Otherwise, the `%dopar%` operator is returned.

Examples

```
`%fun%` <- set_doFuture_strategy("multisession", n_workers=3)
# perform foreach::foreach loop using the %fun% operator
end_doFuture_strategy()
```

 square_loss

Square loss

Description

Square loss

Usage

```
square_loss(y, y_hat)
```

Arguments

y	Vector of observations or ground-truths.
y_hat	Vector of predictions.

Value

The vector of square errors between y and y_hat.

Examples

```
square_loss(c(2.3, 4.2, 1.8), c(2.2, 4.6, 1.7))
```

 unconditional_train_valid_GPD_loss

Unconditional GPD MLEs and their train-validation likelihoods

Description

Unconditional GPD MLEs and their train-validation likelihoods

Usage

```
unconditional_train_valid_GPD_loss(Y_train, interm_lvl, Y_valid)
```

Arguments

Y_train	Vector of "training" observations on which to estimate the MLEs.
interm_lvl	Probability level at which the empirical quantile should be used as the threshold.
Y_valid	Vector of "validation" observations, on which to estimate the out of training sample GPD loss.

Value

Named list containing:

scale	GPD scale MLE inferred from the train set,
shape	GPD shape MLE inferred from the train set,
train_loss	the negative log-likelihoods of the MLEs over the training samples,
valid_loss	the negative log-likelihoods of the MLEs over the validation samples.

vec2mat	<i>Convert a vector to a matrix</i>
---------	-------------------------------------

Description

Convert a vector to a matrix

Usage

```
vec2mat(v, axis = c("col", "row"))
```

Arguments

v	Vector.
axis	One of "col" (default) or "row".

Value

The vector v as a matrix. If axis=="col" (default) the column vector v is returned as a length(v) times 1 matrix. If axis=="row", the vector v is returned as a transposed 1 times length(v) matrix.

Examples

```
vec2mat(c(2, 7, 3, 8), "col")
```

vector_insert	<i>Insert value in vector</i>
---------------	-------------------------------

Description

Insert value in vector

Usage

```
vector_insert(vect, val, ind)
```

Arguments

vect	A 1-D vector.
val	A value to insert in the vector.
ind	The index at which to insert the value in the vector, must be an integer between 1 and $\text{length}(\text{vect}) + 1$.

Value

A 1-D vector of length $\text{length}(\text{vect}) + 1$, with val inserted at position ind in the original vect.

Examples

```
vector_insert(c(2, 7, 3, 8), val=5, ind=3)
```

Index

`%do%`, [25](#), [65](#)
`%dopar%`, [25](#), [65](#)

`backend_is_installed`, [3](#)

`check_directory`, [4](#)
`compute_EQRN_GPDLoss`, [4](#)
`compute_EQRN_seq_GPDLoss`, [5](#)

`default_device`, [6](#)
`default_device()`, [5](#), [6](#), [8](#), [9](#), [11](#), [15–19](#),
[21–23](#), [33](#), [44](#), [45](#), [54](#), [55](#), [60](#), [61](#)

`end_doFuture_strategy`, [7](#)
`EQRN_excess_probability`, [7](#), [20](#)
`EQRN_excess_probability()`, [21](#)
`EQRN_excess_probability_seq`, [8](#), [21](#)
`EQRN_excess_probability_seq()`, [22](#)
`EQRN_fit`, [9](#)
`EQRN_fit()`, [12](#)
`EQRN_fit_restart`, [12](#)
`EQRN_fit_restart()`, [9](#), [13](#)
`EQRN_fit_seq`, [13](#)
`EQRN_fit_seq()`, [12](#), [33](#)
`EQRN_load`, [15](#)
`EQRN_load()`, [19](#)
`EQRN_predict`, [16](#), [44](#)
`EQRN_predict()`, [44](#)
`EQRN_predict_params`, [16](#)
`EQRN_predict_params_seq`, [17](#)
`EQRN_predict_seq`, [18](#), [44](#)
`EQRN_predict_seq()`, [45](#)
`EQRN_save`, [19](#)
`EQRN_save()`, [15](#)
`excess_probability`, [20](#)
`excess_probability.EQRN_iid`, [20](#)
`excess_probability.EQRN_seq`, [21](#)

`FC_GPD_net`, [22](#)
`FC_GPD_SNN`, [23](#)
`fit_GPD_unconditional`, [24](#)

`foreach::foreach()`, [25](#), [65](#)
`future::availableCores()`, [65](#)

`get_doFuture_operator`, [25](#)
`get_doFuture_operator()`, [65](#)
`get_excesses`, [25](#)
`GPD_excess_probability`, [26](#)
`GPD_quantiles`, [27](#)

`install_backend`, [28](#)
`ismev::gpd.fit()`, [24](#)

`lagged_features`, [28](#)
`last_elem`, [29](#)
`loss_GPD`, [29](#)
`loss_GPD_tensor`, [30](#)

`make_folds`, [31](#)
`mean_absolute_error`, [31](#)
`mean_absolute_error()`, [35](#)
`mean_squared_error`, [32](#)
`mean_squared_error()`, [36](#)
`mts_dataset`, [33](#)
`multilevel_exceedance_proba_error`, [34](#)
`multilevel_MAE`, [35](#)
`multilevel_MSE`, [36](#)
`multilevel_pred_bias`, [37](#)
`multilevel_prop_below`, [38](#)
`multilevel_q_loss`, [39](#)
`multilevel_q_pred_error`, [40](#)
`multilevel_R_squared`, [42](#)
`multilevel_resid_var`, [41](#)

`perform_scaling`, [43](#)
`predict.EQRN_iid`, [43](#)
`predict.EQRN_seq`, [44](#)
`predict.QRN_seq`, [45](#)
`predict_GPD_semiconditional`, [47](#)
`predict_unconditional_quantiles`, [48](#)
`prediction_bias`, [46](#)
`prediction_bias()`, [37](#)

prediction_residual_variance, 46
prediction_residual_variance(), 41
process_features, 49
proportion_below, 50
proportion_below(), 38

QRN_fit_multiple, 51
QRN_fit_multiple(), 52
QRN_seq_fit, 52
QRN_seq_fit(), 33, 52, 55, 57
QRN_seq_predict, 45, 54
QRN_seq_predict(), 45
QRN_seq_predict_foldwise, 55
QRN_seq_predict_foldwise(), 56
QRN_seq_predict_foldwise_sep, 56
QRNN_RNN_net, 50
quantile_exceedance_proba_error, 57
quantile_exceedance_proba_error(), 34
quantile_loss, 58
quantile_loss(), 39
quantile_loss_tensor, 59
quantile_prediction_error, 59
quantile_prediction_error(), 40

R_squared, 62
R_squared(), 42
Recurrent_GPD_net, 60
round(), 61
roundm, 61

safe_save_rds, 62
saveRDS(), 62
semiconditional_train_valid_GPD_loss,
63
Separated_GPD_SNN, 64
set_doFuture_strategy, 65
square_loss, 66

torch::dataset, 33
torch::install_torch(), 28
torch::nn_module, 10, 22–24, 50, 51, 60, 61,
64, 65
torch::torch_device(), 5, 6, 8, 9, 11,
15–19, 21–23, 33, 44, 45, 54, 55, 60,
61
torch::torch_is_installed(), 3
torch::torch_save(), 54

unconditional_train_valid_GPD_loss, 66
vec2mat, 67
vector_insert, 68